Tarım Ekonomisi Dergisi
Turkish Journal of Agricultural Economics
ISSN 1303-0183
e-ISSN 3023-5596

Forecasting crop yields with climate and economic variables: A machine learning approach for Türkiye

Burç Arslan KALELİ
Orcid: 0009-0009-1060-1137
Independent Researcher, İstanbul, Türkiye

Makale Künyesi

Araştırma Makalesi / Research Article

Sorumlu Yazar / Corresponding Author
Burç Arslan KALELİ
burckaleli@gmail.com

Geliş Tarihi / Received: 15.05.2025

Kabul Tarihi / Accepted: 01.09.2025

Tarım Ekonomisi Dergisi
Cilt: 31 Sayı: 2 Sayfa: 271-283

Turkish Journal of Agricultural Economics
Volume:31 Issue: 2 Page: 271-283

DOI
10.24181/tarekoder.1699618

JEL Classification: Q10, Q54, C23, C53

Abstract

Purpose: This study forecasts the agricultural productivity of five major crops—wheat, barley, maize, sunflower, and cotton—in Türkiye from 1962 to 2022, using climate variables alone and in combination with economic inputs.

Design/Methodology/Approach: A panel dataset was constructed by matching annual crop yields with seasonal and annual temperature and precipitation variables, including lagged climate indicators. Two model configurations were tested: (i) climate-only and (ii) climate plus economic controls (fertilizer use, capital stock, labor). Three supervised learning models—Linear Regression, Random Forest, and Gradient Boosting—were evaluated using forward-chaining time-series cross-validation.

Findings: Gradient Boosting with economic controls achieved the best out-of-sample performance ($R^2 = 0.44$, MAE = 547.8 kg/ha), followed by Random Forest. Climate-only versions of the same models yielded substantially lower accuracy (e.g., Gradient Boosting $R^2 = 0.16$), highlighting the added predictive value of structural inputs. Feature importance analysis identified growing season temperature as the most influential climate variable, while fertilizer, capital, and labor emerged as key predictors when included.

Originality/Value: This study introduces a robust, time-aware machine learning framework for forecasting crop yields under climate variability. By integrating economic inputs, it enhances predictive accuracy and offers practical insights to support data-driven agricultural planning under climate uncertainty.

Keywords: Agricultural productivity, climate variability, economic inputs, machine learning, yield forecasting.

İklimsel ve ekonomik değişkenlerle ürün verimi tahmini: Türkiye için makine öğrenmesine dayalı bir yaklaşım

Özet

Amaç: Bu çalışma, Türkiye'de buğday, arpa, mısır, ayçiçeği ve pamuk gibi beş temel tarım ürününün yıllık verimleri ile iklim değişkenleri arasındaki tahmine dayalı ilişkiyi 1962–2022 dönemi için incelemektedir. Sadece iklim verilerini kullanan modeller ile iklim verilerine ekonomik girdilerin (gübre kullanımı, sermaye stoku, işgücü) eklendiği modeller karşılaştırılmıştır.

Tasarım/Metodoloji /Yaklaşım: Yıllık ürün verimlerinin, mevsimsel ve yıllık sıcaklık/yağış değişkenleri ile gecikmeli iklim göstergeleriyle eşleştirildiği panel bir veri seti oluşturulmuştur. İki farklı modelleme yapılandırması uygulanmıştır: (i) yalnızca iklim verilerine dayalı ve (ii) iklim + ekonomik kontrol değişkenlerini içeren modeller. Doğrusal regresyon, rastgele orman (Random Forest) ve Gradient Boosting algoritmaları, zaman serisi özellikli ileri zincirleme çapraz doğrulama ile değerlendirilmiştir.

Bulgular: Gradient Boosting modeli, ekonomik kontrol değişkenleriyle birlikte kullanıldığında örneklem dışı performans açısından en başarılı sonuçları vermiştir. ($R^2 = 0.44$, MAE = 547.8 kg/ha). Sadece iklim verileriyle çalışan aynı modelin başarımı belirgin şekilde daha düşüktür (örneğin, $R^2 = 0.16$). Özellik önem analizleri, büyüme dönemindeki sıcaklığın iklim değişkenleri arasında en belirleyici unsur olduğunu; gübre, sermaye ve işgücünün ise eklendiğinde tahmin gücünü önemli ölçüde artırdığını göstermektedir.

Özgünlük/Değer: Bu çalışma, iklim değişkenlerine ekonomik yapısal girdileri entegre eden, zamana duyarlı ve güçlü bir makine öğrenmesi çerçevesi sunmayı amaçlamaktadır. Bu sayede tarımsal üretimin iklimsel belirsizliklere karşı hassasiyeti daha doğru tahmin edilebilmekte ve veriye dayalı tarımsal planlama süreçlerine katkı sağlanmaktadır.

Anahtar kelimeler: Tarımsal verimlilik, iklim değişkenliği, ekonomik girdiler, makine öğrenmesi, verim tahmini.

INTRODUCTION

Climate change is one of the most critical challenges of the 21st century, with far-reaching consequences for food security, rural livelihoods, and macroeconomic stability. These effects are particularly acute in regions with climate-sensitive agricultural systems and limited adaptive capacity. The Mediterranean basin is considered a climate change “hotspot” due to its susceptibility to long-term droughts, high temperatures, and erratic precipitation patterns (Cramer et al., 2018). Türkiye, with its semiarid geography and reliance on agriculture, is especially vulnerable. Agriculture remains a major socioeconomic pillar in Türkiye, employing 17.1% of the national workforce in 2022 (TÜİK, 2023) and representing the primary livelihood in many rural areas.

Due to widespread dependence on rainfed agriculture and seasonal precipitation, Türkiye's cropping systems are highly sensitive to climatic fluctuations (Bakırcı and Çakır, 2023). Key crops such as wheat, barley, maize, sunflower, and cotton dominate both harvested area and calorie intake (FAO, 2023; TÜİK, 2023) and are especially exposed to heat and water stress (Özbek and Özbek, 2024). Model projections suggest that climate change could reduce cereal yields in the Mediterranean by 10–25% under moderate emissions scenarios (IPCC, 2021; Zampieri et al., 2017). National studies indicate that intensifying temperature trends, soil degradation, and precipitation variability already pose significant risks to productivity in Türkiye (Yaraşır et al., 2023; Ortaş, 2024).

While climate-yield relationships have been widely studied, the literature reflects varying perspectives on which methodologies best capture these dynamics. Traditional econometric models—particularly fixed-effect panel regressions—offer interpretability and the ability to control for unobserved heterogeneity (Deschênes and Greenstone, 2007), but they are often limited in addressing nonlinear, threshold-based, or interacting climate effects common in biological systems.

Advances in data science and machine learning (ML) provide alternative tools for modeling complex, nonlinear relationships. Ensemble models such as Random Forest and Gradient Boosting are particularly suited to handling multivariate interactions and flexible patterns (Kamilaris and Prenafeta-Boldú, 2018; van Klompenburg et al., 2020). However, their limited transparency can pose challenges for use in policy contexts that require interpretability.

In this study, we adopt a machine learning–oriented approach to forecast crop yields in response to climatic and economic variation in Türkiye over the 1962–2022 period. Rather than attempting causal inference, we focus on evaluating the out-of-sample predictive performance of different model configurations and exploring their potential to support agricultural decision-making. We compare Random Forest and Gradient Boosting to a baseline Linear Regression model in two configurations: one using only climate variables, and the other incorporating economic inputs such as fertilizer, capital, and labor.

The models are evaluated using a time-aware cross-validation procedure designed to avoid temporal leakage and better reflect real-world forecasting settings. To improve interpretability, we apply feature importance rankings to identify influential predictors—offering indicative insights that may inform agricultural risk management.

By analyzing long-term climate-yield patterns through this machine learning framework, the study seeks to provide additional empirical evidence on the predictive capacity of alternative model structures and highlight crop-level vulnerabilities that may inform early-warning strategies for agricultural planning in Türkiye.

LITERATURE REVIEW

The relationship between climate change and agricultural productivity is a central concern in environmental economics and agricultural modeling. Within this expanding field, two commonly applied methodological approaches have emerged: econometric modeling and machine learning (ML). While econometrics emphasizes causal inference and theoretical rigor, ML prioritizes predictive accuracy and flexibility in modeling complex, nonlinear systems. This study adopts a purely machine learning–based approach, focusing on yield forecasting rather than causal estimation.

Econometric models have traditionally been used to estimate the influence of climate on crop production, particularly through fixed-effects panel regressions that control for unobserved heterogeneity across space and time. A pioneering example is Deschênes and Greenstone (2007), who exploited weather variability to estimate climate effects on U.S. agriculture. Subsequent studies extended this approach in different agroecological contexts, incorporating variables such as lagged weather, soil quality, and farm-level data to improve robustness.

In Türkiye, several studies have explored regional climate-yield relationships using classical econometric frameworks. Altinsoy et al. (2012) estimated the long-term effects of temperature and precipitation on wheat, barley, and maize across agroclimatic zones. Açci et al. (2024) used a hybrid econometric-ML framework to assess climate-induced food insecurity, highlighting macroeconomic risks stemming from agricultural volatility.

However, limitations of classical models-especially their reliance on linear functional forms—have created opportunities for ML-based alternatives, particularly in forecasting applications. Traditional models may face challenges in capturing nonlinearities, interaction effects, and thresholds that are common in climate and biological systems. ML methods offer greater flexibility in modeling such dynamics. Kamilaris and Prenafeta-Boldú (2018), in a comprehensive review, emphasized that ensemble algorithms such as Random Forest and Gradient Boosting are especially effective at detecting complex relationships in large agronomic datasets.

A meta-analysis by Van Klompenburg et al. (2020), reviewing 50 ML-based crop yield forecasting studies, found that ensemble methods generally outperformed linear regressions, particularly when combining meteorological and agronomic variables. Temperature and precipitation were dominant predictors in over 90% of the reviewed models, reinforcing their empirical significance. The study also discussed trade-offs between algorithms: while Random Forest is computationally efficient, Gradient Boosting tends to deliver higher accuracy at greater computational cost.

A recurring concern in this literature is model interpretability. High-performing ML models can lack transparency, which may limit their usefulness in policy contexts. Several studies have adopted explainability tools to address this issue. Shahhosseini et al. (2020) applied SHAP (SHapley Additive exPlanations) to Gradient Boosting models of wheat yield in Australia and identified maximum reproductive temperature as the leading predictor, offering insight into biological thresholds. Kuwata and Shibasaki (2015) also demonstrated the forecasting power of ensemble models across agroclimatic zones, though interpretability remained limited.

Empirical studies further underscore the predictive advantages of nonlinear methods. Jeong et al. (2016) found that Random Forest consistently outperformed linear models in predicting global and regional yields, particularly under variable spatial inputs. Khaki and Wang (2019) used convolutional neural networks and genetic algorithms to forecast corn yields in the U.S. Corn Belt, outperforming traditional methods across spatial resolutions. Osman et al. (2025) applied a combination of Random Forest and Bayesian methods to forecast climate-related agricultural shocks in Somalia, effectively capturing seasonal thresholds. In Türkiye, Bozdağ (2021) used artificial neural networks to simulate carbon emissions under changing land use and climate scenarios, illustrating the spatial forecasting potential of ML tools for national adaptation planning.

Despite this progress, key gaps remain in Türkiye's literature. Most existing studies are crop-specific or limited to short time periods, reducing generalizability. Few have applied long-term ML-based forecasting frameworks across multiple crops within a unified architecture. This is particularly relevant given Türkiye's agroclimatic variability, reliance on rainfed agriculture, and growing exposure to seasonal temperature and precipitation changes. The present study aims to contribute to this gap by offering a long-term, multi-crop ML-based forecasting framework with an emphasis on predictive performance, explainability, and temporal generalization.

METHODOLOGY

This study employs three supervised learning algorithms to forecast crop yields across five key crops in Türkiye from 1962 to 2022:

(i) Linear Regression (as a baseline comparator),

(ii) Random Forest Regressor, and

(iii) Gradient Boosting Regressor.

Two model configurations were designed to assess the relative predictive contributions of different feature sets:

- Climate-Based Model: Includes only climate features, such as annual and growing season precipitation and temperature, along with one-year lags. Crop types were represented using one-hot encoded indicators to account for differences across crops.

• Extended Feature Model: Augments the climate-based model with structural inputs including fertilizer application (kg/ha), agricultural labor (1,000 persons), and capital stock (million USD). These features were included to improve forecasting performance by incorporating information on resource use and production capacity.

Each model was trained and evaluated separately to isolate the predictive contribution of climate variables versus extended structural features.

Linear Regression

Linear Regression was used as a baseline model. Unlike fixed-effect econometric formulations, this version does not incorporate crop-specific fixed effects or robust standard errors and is applied strictly for predictive purposes. The general structure of the linear regression follows the form presented in Equation (1):

$$\widehat{y_i} = \beta_0 + \sum_{j=1}^{p} \beta_j x_{ij} \quad (1)$$

Where:

- $\widehat{y_i}$ is the predicted yield for observation i,
- x_{ij} are the input features (e.g., climate variables),
- β_j are the estimated coefficients,
- β_0 is the intercept — the value of $\widehat{y_i}$ when all $x_{ij} = 0$
- p is the total number of predictors

While linear regression assumes additive and linear relationships, it serves here as a benchmark to evaluate whether more flexible models provide predictive improvement.

Random Forest Regressor

The Random Forest algorithm is an ensemble method that constructs multiple decision trees on bootstrapped subsets of the data and averages their predictions to reduce overfitting and variance (Breiman, 2001). The decision of the Random Forest model is given in Equation (2):

$$\widehat{f_{RF}}(x) = \frac{1}{B}\sum_{b=1}^{B} T_b(x) \quad (2)$$

Where:

- B is the number of decision trees,
- $T_b(x)$ is the prediction from the $b-th$ tree,
- $\widehat{f_{RF}}(x)$ is the final predicted yield.

Random Forests can model nonlinear relationships and interactions, and they generate feature importance scores based on mean impurity reduction, providing insight into the relative contribution of input variables for prediction.

Gradient Boosting Regressor

Gradient Boosting is a sequential ensemble method in which each tree is trained to correct the residual errors of the prior model (Friedman, 2001). The model is trained iteratively as Equation (3):

$$\widehat{f_m}(x) = \widehat{f_{m-1}}(x) + \nu \cdot h_m(x) \quad (3)$$

Where:

- $\widehat{f_m}(x)$ is the ensemble prediction after m steps,

- $h_m(x)$ is the weak learner (a decision tree) fitted to the residuals,
- v is the learning rate (typically $0 < v \leq 1$).

The advantage of Gradient Boosting over Random Forest is that it generally gives better (higher accuracy) results but is more sensitive to overfitting and requires fine tuning of tree depth, learning rate, and number of predictors.

Performance Metrics

All models are assessed using Mean Absolute Error (MAE), Root Mean Squared Error (RMSE) and Coefficient of Determination (R^2) given in Equation (4):

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|$$
$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \quad (4)$$
$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2}$$

Where:

- y_i is the observed yield,
- $\hat{y}_i$ is the predicted yield,
- $\bar{y}$ is the sample mean of observed yields,
- n is the number of observations.

DATA

Data description

This study uses a panel dataset covering crop yields, climate conditions, and economic inputs for Türkiye over the 1962–2022 period. Data were compiled from FAOSTAT, the World Bank Climate Portal, and the USDA Economic Research Service (as of October 2024). Yield data (kg/ha) were collected for five key crops: wheat, barley, maize, sunflower seed, and seed cotton—selected for their national importance and climate sensitivity.

Climate data originate from the CRU TS v4.07 dataset, containing monthly precipitation and surface temperature values. These were aggregated to both annual and growing season (April–September) levels to reflect critical crop development periods. To account for possible lagged effects—such as residual soil moisture or heat accumulation-one-year lagged values of annual temperature and precipitation were also included. This enables the model to capture temporal dependencies relevant for semi-perennial systems and cumulative climate stress.

Three economic input variables were also incorporated to enhance model accuracy and generalizability:

- Fertilizer (kg/ha): Calculated as the ratio of total fertilizer use (in metric tons of inorganic nitrogen, phosphorus, potassium, and organic nitrogen) to cropland area (in 1,000 hectares).
- Capital (million $): Defined as the total agricultural capital stock at constant 2015 prices.
- Labor (1,000 persons): defined as the total agricultural labor force economically active in farming.

These features were collected annually and uniformly applied across crop-year pairs. Crop-specific effects were handled through one-hot encoding. The resulting panel dataset consists of 305 observations (5 crops × 61 years), structured to support both temporal forecasting and crop-specific modeling.

Although this study does not pursue causal estimation, the inclusion of economic variables supports more robust predictions by reflecting structural patterns in agricultural inputs. Table 1 summarizes variable definitions.

Table 1. Description of variables

Variable	Type	Construction	Notes
year	Time index	Extracted from all files	Used for merging and alignment
crop	Categorical	FAOSTAT	Wheat, Barley, Maize, Sunflower, Cotton
yield_kg_ha	Target variable	FAOSTAT (kg/ha)	Yield per hectare
precip_mm	Climate Input	Total monthly rainfall, summed per year	Annual precipitation
temp_c	Climate Input	Average of monthly temperature per year	Annual mean temperature
rainfall_q2_q3	Climate Input	Sum of rainfall April–September	Growing season precipitation
temperature_q2_q3	Climate Input	Avg. temperature April–September	Growing season temperature
rainfall_lag1	Lagged Climate Input	Lag of precip_mm (t−1)	Captures delayed soil moisture effects
temperature_lag1	Lagged Climate Input	Lag of temp_c (t−1)	Reflects thermal memory
fertilizer	Economic Input	Fertilizer use / cropland area	kg/ha; Source: USDA ERS
Capital	Economic Input	Total agricultural capital stock	$million, constant 2015 prices
Labor	Economic Input	Total agricultural labor	In 1,000 persons

To facilitate comparability across years and crops, all variables were harmonized and merged into a unified structure. This enables the machine learning models to leverage both interannual and cross-crop variation.

Table 2 summarizes the descriptive statistics of key variables. Crop yields show wide variability (mean = 2,633 kg/ha; SD = 1,754), consistent with Türkiye's agroecological diversity. Climate indicators vary meaningfully: annual rainfall ranges from 444 mm to 738 mm, and annual temperature from 9.7°C to 13.1°C. Growing season values average 208 mm for precipitation and 17.8°C for temperature—highlighting the seasonal climatic exposure faced by crops.

Among economic inputs, fertilizer use ranges from 5 to 131 kg/ha. Capital stock expanded from $31 billion to $133 billion during the study period, while labor input averaged 7.2 million persons, declining gradually over time—reflecting structural transformation in the agricultural sector.

Table 1. Summary statistics of key variables (1962–2022)

Variable	Mean	Std Dev	Min	Q1	Median	Q3	Max
Yield (kg/ha)	2632.9	1753.84	738	1548.5	2122.7	2786	9635.8
Annual Rainfall (mm)	598.52	69.87	443.66	545.18	608.28	641.3	738.15
Annual Temp (°C)	11.3	0.74	9.69	10.72	11.29	11.79	13.06
Q2–Q3 Rainfall (mm)	208.06	36.75	120.1	192.19	204.8	226	312.04
Q2–Q3 Temp (°C)	17.82	0.73	16.39	17.22	17.73	18.46	19.2
Lagged Rainfall (mm)	599.02	69.48	443.66	550.53	608.28	641.3	738.15
Lagged Temp (°C)	11.29	0.73	9.69	10.72	11.27	11.78	13.06
Fertilizer (kg/ha)	63.01	31.74	4.98	47.2	65.69	83.12	130.71
Capital ($million)	69300	27944.1	31102	50252	60704.3	85448	132902
Labor (1,000 persons)	7255.4	1432.6	4705.2	5713	7808.4	8207	9259

Figure 1a displays yield distributions by crop. Maize and cotton exhibit the highest median yields and the widest variability, while barley and wheat show lower medians with tighter interquartile ranges. These differences support the inclusion of crop-level identifiers in stratified modeling to capture heterogeneity in crop response.

Figures 1b–1d present long-term changes in agricultural and climatic indicators. Average yield (Figure 1b) increased markedly from 1980 to 2010, with a plateauing trend thereafter that may reflect broader systemic or climatic constraints. Annual precipitation (Figure 1c) fluctuates without a clear trend, reinforcing the decision to emphasize seasonal indicators. Average annual temperature (Figure 1d) shows a steady rise—from approximately 10.5°C to 12.5°C—consistent with warming patterns across the region, albeit with considerable inter-annual variation.

Economic input variables display distinct temporal dynamics. Capital stock (Figure 1e) shows a consistent upward trend, accelerating notably after 2000, in line with long-term investment patterns. Fertilizer use (Figure 1f) increased rapidly until the 1980s and has continued to rise overall, albeit with notable year-to-year fluctuations likely driven by policy and market factors. These evolving trends justify the inclusion of structural economic controls in the model to capture broader shifts in input intensity over time.

Figure 1. Descriptive summary of crop, climate data, and economic inputs in türkiye 1962–2022 – (a) Yield Distributions by Crop Type (b) Average Crop Yields (c) Total Annual Precipitation (d) Average Annual Temperature (e) Agricultural Capital Stock (f) Fertilizer Used per Hectare

Data preprocessing

To prevent temporal leakage and ensure realistic model evaluation, the dataset was split chronologically using the TimeSeriesSplit method from Scikit-learn. This approach strictly preserves temporal order by training models on past data and testing on future observations. Such a design is essential for yield forecasting, where future conditions cannot be known in advance.

A five-fold, time-aware cross-validation scheme was employed. In each fold, the training window expands forward in time, while the test window moves ahead by approximately a decade. The structure of these folds is summarized in Table 3.

Table 3. Time-aware cross-validation structure (forward-chaining 5-fold split)

Fold	Training Period	Testing Period
1	1962–1972	1973–1982
2	1962–1982	1983–1992
3	1962–1992	1993–2002
4	1962–2002	2003–2012
5	1962–2012	2013–2022

Each observation retains its crop identity through one-hot encoding of the crop variable. This allows the model to learn crop-specific yield dynamics without introducing interference across categories.

The fold structure is visualized in Figure 2, with training periods shaded in blue and testing periods in orange, clearly illustrating the forward-chaining progression.

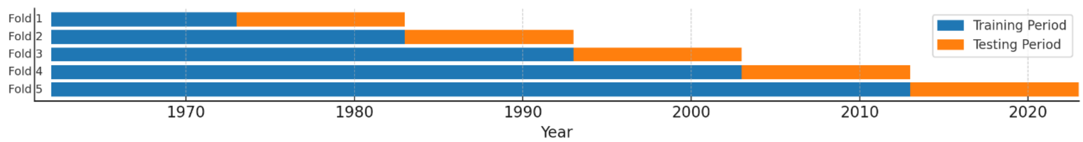

Figure 2. Structure of time-based folds used in cross-validation

APPLICATIONS and RESULTS

Model configurations and training procedure

To evaluate predictive performance under different modeling strategies, three supervised learning algorithms—Linear Regression, Random Forest, and Gradient Boosting—were applied to two configurations: one using climate variables only, and another combining climate with structural economic inputs. Each model was trained under a five-fold, forward-chaining cross-validation scheme, with training windows expanding over time and all testing periods confined to future years. This setup replicates real-world forecasting and prevents temporal leakage.

Performance metrics (MAE, RMSE, R^2) were calculated on out-of-sample test folds and averaged. The following results section compares all model–configuration combinations and discusses predictive drivers, feature importance analysis and crop-specific effects.

Performance comparison across models and specifications

The predictive performance of each model–specification combination is summarized in Table 4 and visualized in Figure 3. Across all metrics—Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R^2—Gradient Boosting with economic inputs yielded the most accurate forecasts. This model achieved an MAE of approximately 547.8 kg/ha and an RMSE of 812.5 kg/ha, with an R^2 of 0.44, indicating that about 44% of the variation in yield was captured in out-of-sample forecasts.

Table 4. Model performance comparison across configurations

Model	Controls	MAE ((kg/ha))	RMSE ((kg/ha))	R^2
Linear Regression	With Econ	887.7	1160.5	0.32
Random Forest	With Econ	569.9	829.51	0.43
Gradient Boosting	With Econ	547.8	812.47	0.44
Linear Regression	Climate Only	956.7	1331.7	0.05
Random Forest	Climate Only	904.2	1183.3	0.15
Gradient Boosting	Climate Only	902.7	1193.2	0.16

In contrast, the same model applied without economic inputs (Climate Only) showed a notable drop in performance, with a higher MAE of 902.7 and a lower R^2 of 0.16. The improved accuracy when including fertilizer, capital, and labor suggests these structural variables provide meaningful predictive information, especially in long-term forecasts.

Among the three models, Linear Regression consistently performed the weakest. Although its R^2 increased from 0.05 to 0.32 with the addition of economic controls, its MAE and RMSE remained significantly higher than those of the ensemble models. This reflects the linear model's limitations in capturing nonlinearities and interactions common in crop-climate relationships.

Both Random Forest and Gradient Boosting improved substantially when economic inputs were added. Gradient Boosting slightly outperformed Random Forest in both configurations, particularly in terms of RMSE and explained variance. These gains likely reflect Gradient Boosting's iterative learning and stronger bias reduction capability.

Figure 3 further illustrates these patterns. All models perform worse when limited to climate variables alone, with R^2 dropping sharply in all cases. The most pronounced improvement is seen in the Gradient Boosting model, where the inclusion of economic inputs raises R^2 by over 28 percentage points. This consistent gain across all models confirms that economic structural factors materially enhance the predictive capacity of climate-yield models.

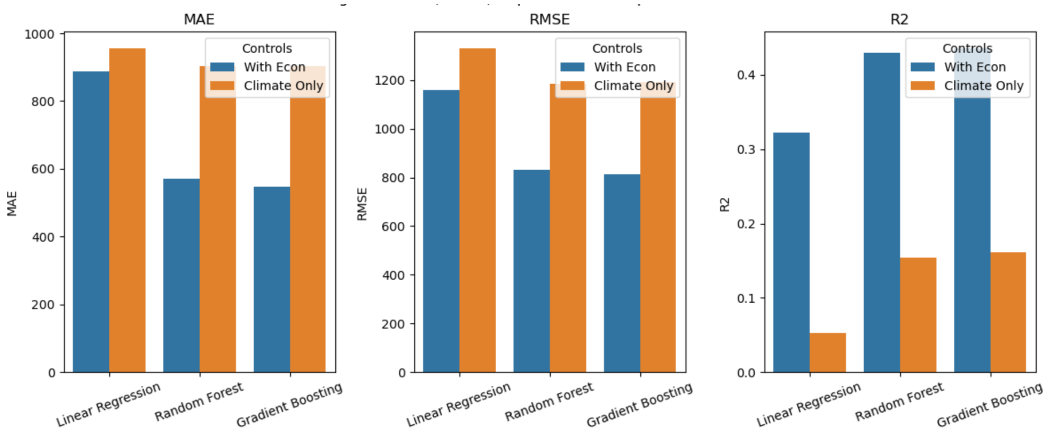

Figure 3. MAE, RMSE, and R^2 by model and specification

Results underscore the value of integrating economic and climatic drivers within a machine learning framework. They also highlight the critical role of model specification, with Gradient Boosting emerging as the most effective across both feature sets.

Feature importance analysis

To better understand the predictive mechanisms behind model performance, we examined feature importances from the best-performing model—Gradient Boosting—under both configurations: with and without economic inputs. These importances were averaged across all five test folds and then separately recalculated using only the final fold to evaluate temporal consistency and any late-period shifts in driver relevance.

Figure 4 presents the top features across the full validation scheme. In the climate-only model (right panel), temperature during the growing season (temperature_q2_q3) emerged as the single most important predictor of yield, followed by lagged and seasonal rainfall variables. These findings align with agronomic expectations, as temperature directly influences photosynthesis and crop development phases.

However, once structural economic inputs were introduced (left panel), the feature landscape shifted markedly. Fertilizer input, capital stock, and labor rose to the top, displacing climatic variables in relative importance. Among these, fertilizer exhibited the highest predictive contribution, followed by capital and labor, indicating the pronounced role of input intensification and mechanization in yield variation. Importantly, this re-ranking does not negate the role of climate but suggests that economic heterogeneity adds substantial predictive power when modeled jointly.

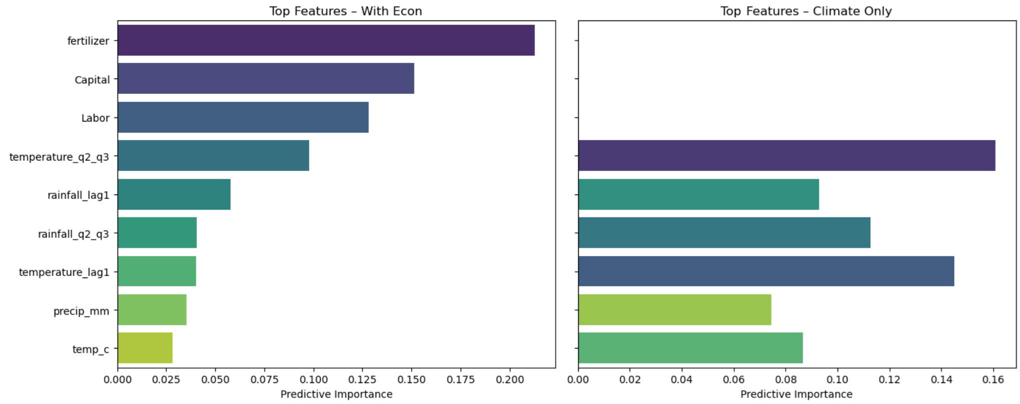

Figure 4. Feature importances in gradient boosting – five-fold average

To probe the temporal stability of these patterns, Figure 5 isolates feature weights derived from the final out-of-sample window. In this late period (2013–2022), capital stock emerged as the top-ranked feature, closely followed by labor and fertilizer. This reordering—compared to the five-fold average—may reflect recent shifts in production technology, mechanization levels, or structural investment. Notably, temperature_q2_q3 maintained a strong role as the leading climate variable in both settings, reinforcing its fundamental biological relevance.

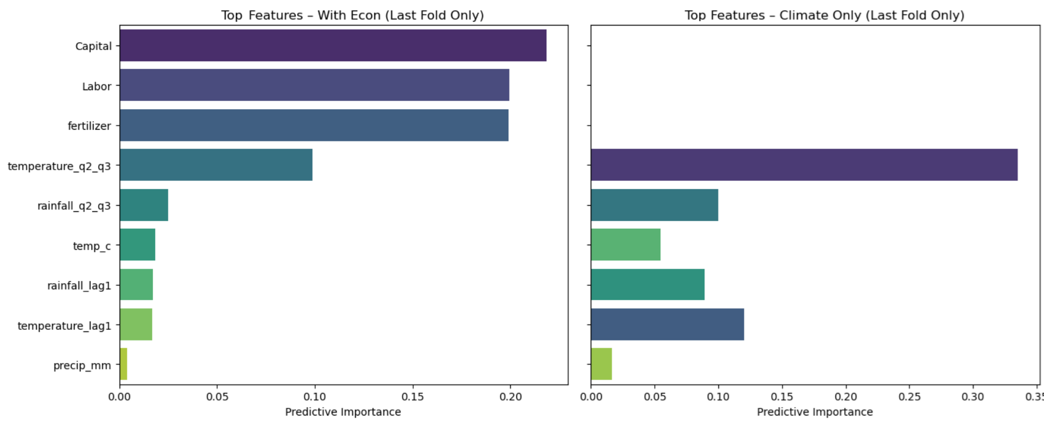

Figure 5. Importances in Gradient Boosting – Last Fold Only

Together, these results indicate that the observed improvement in predictive performance (Section 5.2) stems not from one or two dominant features, but from the inclusion of a meaningful and temporally stable set of structural agricultural factors. The consistency of economic variables across both multi-fold and recent-year analyses suggests they are not only correlated with yield but contribute systematically to its prediction over time. While these models do not claim causal identification, the relative importance rankings provide transparent insight into how different features influence yield forecasts under the chosen learning algorithms.

Crop-specific model behavior

To investigate whether the observed performance improvements were consistent across different crops, we analyzed the distribution of absolute prediction errors by crop and model configuration using the Gradient Boosting algorithm. Figure 6 presents boxplots of mean absolute error (MAE) across all five test folds for each crop, comparing the climate-only model (orange) with the version incorporating economic variables (blue).

The results reveal notable heterogeneity in both baseline error magnitudes and the degree of improvement achieved through economic enrichment. For all crops, the addition of structural controls reduced prediction errors, though the extent varied by crop. The most substantial gains were observed in high-variability, input-sensitive crops such as maize (corn) and seed cotton, where both the median and upper-tail errors declined sharply. For instance, in maize, the interquartile range and median MAE dropped visibly, suggesting that input-related heterogeneity—such as fertilizer responsiveness or mechanization intensity—plays a significant role in yield volatility that purely climatic factors cannot fully explain.

In contrast, crops like barley and wheat exhibited lower overall prediction errors in both configurations, and the marginal improvement from economic variables was correspondingly smaller. This may reflect their greater resilience to input variation or more stable production technologies. Nevertheless, even in these cases, the economic model consistently achieved tighter error distributions and lower median MAEs, implying more stable forecasts across years.

The results for sunflower seed further support this interpretation. Despite exhibiting relatively moderate error levels, the addition of economic inputs resulted in both a downward shift in the error distribution and a visible compression of outliers. This suggests that while climatic conditions remain important, structural production variables help moderate extreme forecast deviations—particularly for crops with more elastic input-output relationships.

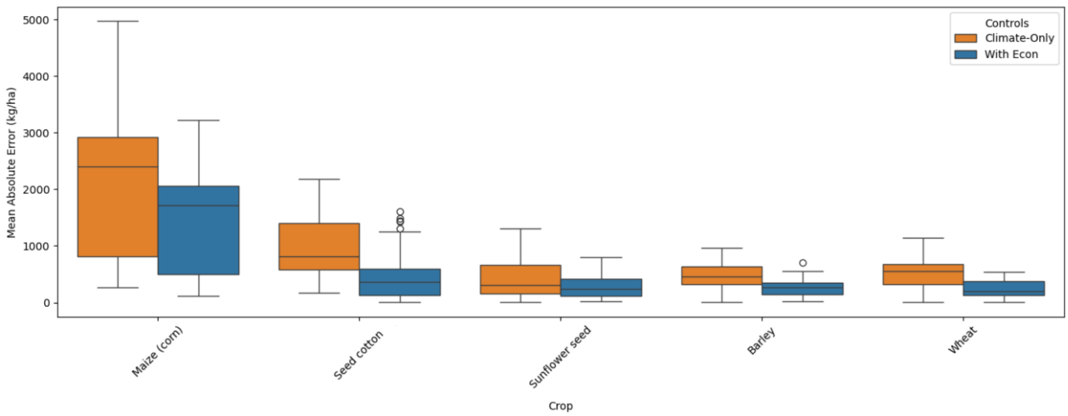

Figure 6. Crop-specific MAE distribution (all folds, gradient boosting)

Overall, these crop-level patterns reinforce the conclusion that incorporating economic variables improves not only average predictive performance but also cross-sectional consistency. The gains are especially pronounced for crops that are more sensitive to farm-level investment decisions and input intensity. This heterogeneity underlines the importance of avoiding a one-size-fits-all modeling approach when forecasting yields across diverse agricultural systems.

Interpretation and policy relevance

The predictive findings of this study offer several insights for agricultural planning and adaptive decision-making. Among the tested approaches, the Gradient Boosting model—particularly when supplemented with structural variables—consistently yielded superior out-of-sample performance. This suggests that integrating economic inputs such as fertilizer use, capital stock, and labor availability with climatic indicators provides a more complete basis for anticipating yield variability in the face of uncertain growing conditions.

Temperature during the second and third quarters emerged as the most influential climate feature in the absence of economic data. However, once structural variables were introduced, the predictive hierarchy shifted, with fertilizer, capital, and labor rising in relative importance. This reordering reflects predictive salience, not causal effect. These results do not imply that economic inputs drive yield outcomes independently of climate, but rather that they carry useful information about underlying production dynamics when included in forecasting frameworks.

From a policy relevance standpoint, the differentiated behavior across crops implies that a uniform intervention strategy may not be optimal. For crops like maize and cotton, where prediction accuracy improved most with the addition of economic controls, support policies might focus on enhancing access to fertilizers, improving credit availability, or promoting mechanization. In contrast, barley and wheat, which showed more stable yield patterns and lower baseline errors, may benefit more from climate monitoring tools, drought alerts, or extension services focused on weather-risk mitigation.

These applications must be interpreted cautiously. The models are explicitly predictive tools, optimized for forecasting accuracy—not explanatory inference. As such, variable importance should not be taken as evidence of structural relationships or as a basis for policy evaluation.

Moreover, while this study includes key structural indicators, it does not incorporate farm-level variables such as irrigation access, soil characteristics, or market linkages, which may influence both yield variability and model interpretability. Similarly, national-level averages likely mask subregional climate-agriculture dynamics, which are especially salient in a geographically diverse country like Türkiye.

Nonetheless, the observed performance gains suggest that machine learning models, when transparently constructed and properly framed, can play a valuable role in short- to medium-term agricultural forecasting. They are particularly useful when used in conjunction with improved data systems and expert domain knowledge. The results presented here are intended to complement—not replace—causal analyses, and may support the development of risk-aware planning tools, adaptive input allocation policies, and early warning systems for agriculture.

CONCLUSION

This study examined the predictive relationship between climate conditions, structural economic inputs and agricultural productivity in Türkiye over the 1962–2022 period, focusing on five major crops: wheat, barley, maize, sunflower, and cotton. Using a strictly machine learning–based approach, we trained and validated three supervised models—Linear Regression, Random Forest, and Gradient Boosting—across both climate-only and extended feature sets that included structural economic variables. To preserve the integrity of the temporal sequence and prevent forward-looking bias, all models were evaluated using a forward-chaining cross-validation strategy.

The results highlight the value of incorporating both climatic and structural economic inputs into forecasting models. In particular, growing season temperature consistently appeared among the most influential predictors, while the addition of economic variables—fertilizer, capital stock, and labor—significantly improved model stability and reduced error across all crops. These improvements were especially evident for input-sensitive crops such as maize and cotton, reinforcing the importance of modeling heterogeneity in production systems.

Throughout this analysis, care was taken to maintain a clear distinction between predictive modeling and causal inference. The goal was not to estimate policy-invariant parameters or test theoretical relationships, but rather to develop models that could reliably anticipate yield outcomes based on historical input patterns.

While the findings provide actionable insights for yield forecasting and early warning applications, several limitations must be noted. The use of national-level averages likely conceals significant subregional variation, particularly relevant in a country with Türkiye's geographic and climatic diversity. Furthermore, key factors such as irrigation availability, soil quality, and farm-level management practices were not included due to data constraints, which may limit the full explanatory power of the models. The limited predictive role of lagged climate variables suggests that either longer-term memory effects are weak or that annual data do not capture such dynamics with sufficient resolution.

Looking ahead, future work would benefit from incorporating spatially disaggregated datasets and exploring the influence of extreme weather events on model performance. Capturing dynamic adaptation behaviors, including multi-year responses to persistent shocks or investment shifts, may also enhance forecasting accuracy. Machine learning models such as those used here are not substitutes for structural economic analysis but can serve as complementary tools—especially when high-frequency forecasts and operational insights are needed.

As a result, this study demonstrates that predictive models, when carefully constructed and transparently interpreted, offer a practical contribution to agricultural risk management. By integrating climate and economic variables within a time-aware machine learning framework, the results provide a basis for data-driven decision-making in agricultural planning and support the development of adaptive strategies under changing environmental conditions.

REFERENCES

Açci, Y., Uçar, E., Uçar, M. and Açci, R.C. (2024), "Evaluating the relationship between climate change, food prices, and poverty: empirical evidence from underdeveloped countries", *Environment Development and Sustainability.*

Altinsoy, H., Kurt, C. and Kurnaz, M.L. (2012), "Analysis of the effect of climate change on the yield of crops in Turkey using a statistical approach", *Springer Atmospheric Sciences*, pp. 379–384.

Bakırcı, G.T. and Çakır, Ö. (2023), "Evaluation of climate change in the scope of agriculture and food", Akademik Gıda, 2024, pp. 57-64.

Bozdağ, A. (2021), "Local-based mapping of carbon footprint variation in Turkey using artificial neural networks", Arabian Journal of Geosciences, 14(6), pp.486.

Breiman, L. (2001), "Random forests", *Machine Learning*, Vol. 45 No. 1, pp.5–32.

Cramer, W., Guiot, J., Fader, M., Garrabou, J., Gattuso, J.-P., Iglesias, A., Lange, M.A., Lionello, P., Llasat, M.C., Paz, S. et al. (2018), "Climate change and interconnected risks to sustainable development in the Mediterranean", Nature Climate Change, 8(11), pp.972–980.

Deschênes, O. and Greenstone, M. (2007), "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather", *American Economic Review*, Vol. 97 No. 1, pp. 354–385.

FAO (2023) FAOSTAT: Turkey – Crops and Livestock Products, Food and Agriculture Organization of the United Nations, Rome. Available at: https://www.fao.org/faostat/en/#country/223 (Accessed: 29 April 2025).

Friedman, J.H. (2001), "Greedy function approximation: A gradient boosting machine", Annals of Statistics, 29(5), pp.1189–1232.

IPCC (2021), Climate Change 2021: The Physical Science Basis, Intergovernmental Panel on Climate Change, Geneva. Available at: https://www.ipcc.ch/report/ar6/wg1/ (Accessed: 1 April 2025).

Jeong, J.H., Resop, J.P., Mueller, N.D., Fleisher, D.H., Yun, K., Butler, E.E., Timlin, D.J., *et al.* (2016), "Random forests for global and regional crop yield predictions", *PLoS ONE*, Vol. 11 No. 6, p. e0156571.

Kamilaris, A. and Prenafeta-Boldú, F.X. (2018), "Deep learning in agriculture: A survey", *Computers and Electronics in Agriculture*, Vol. 147, pp. 70–90.

Khaki, S. and Wang, L. (2019), "Crop yield prediction using deep neural networks", *Frontiers in Plant Science*, Vol. 10.

Kuwata, K. and Shibasaki, R. (2015), "Estimating Crop Yields with Deep Learning and Remotely Sensed Data", *Proceedings of the 2015 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 858–861.

Ortaş, İ. (2024), "Under Long-Term Agricultural Systems, the role of mycorrhizae in climate change and food security", *Manas Journal of Agriculture Veterinary and Life Sciences*, 14(1), pp. 101–115.

Osman, B.M., Mohamud, M.H., Omar, A.F., Yabarow, A.A., Omar, O.M., Jirow, M.S.A., Dahir, A.N., *et al.* (2025), "The impact of climate change on economic growth in Somalia: using random forest and Bayesian approach", *Cogent Economics & Finance*, 13(1).

Özbek, S. and Özbek, B. (2024), "Does Climate Change Strengthen the Link between Environmental Degradation and Agricultural Output? Empirical Evidence on the Turkish Economy", *Tarım Ekonomisi Dergisi*, 30(1) pp. 49–60.

Rossini, A., Ruggeri, R., Belocchi, A. and Rossini, F. (2024), "Response of durum wheat cultivars to climate change in a Mediterranean environment: Trends of weather and crop variables at the turn of 21st century", *Journal of Agronomy and Crop Science*, 210(6).

Shahhosseini, M., Hu, G. and Archontoulis, S.V. (2020), "Forecasting corn yield with machine learning ensembles", *Frontiers in Plant Science*, Vol. 11.

TÜİK (2023), *Labour Force Statistics 2022*, Turkish Statistical Institute, Ankara. Available at: https://data.tuik.gov.tr/Bulten/Index?p=Labor-Force-Statistics-2022-49167 (Accessed: 25 April 2025).

Van Klompenburg, T., Kassahun, A. and Catal, C. (2020), "Crop yield prediction using machine learning: A systematic literature review", *Computers and Electronics in Agriculture*, Vol. 177, p. 105709.

Yaraşır, N., Yiğit, A. and Erekul, O. (2024), "Yield and quality responses of soybean (Glycine max L. Merr.) varieties inoculated with rhizobia strains under drought stress", *Turkish Journal of Field Crops*, 29(2), 165-176.

Zampieri, M., Ceglar, A., Dentener, F. and Toreti, A. (2017), "Wheat yield loss attributable to heat waves, drought and water excess at the global, national and subnational scales", *Environmental Research Letters*, Vol. 12 No. 6, p. 064008.